Michael Blaugrund
Head of Transactions

New York Stock Exchange
11 Wall Street
New York, NY 10005
T: [redacted]

October 24, 2018

VIA E-MAIL

Brent J. Field
Secretary
Securities and Exchange Commission
100 F Street, N.E.
Washington, DC 20549-1090

Re: Comments for Consideration for Panel 4 of the SEC's Roundtable on Market Data and Market Access (File No. 4-279)

Dear Mr. Fields:

NYSE Group is submitting this comment letter to provide additional information in connection with my participation in Panel 4 - Elements of the Core Data Infrastructure relating to ways to modernize the infrastructure underlying the receipt, processing, and delivery of core market data by the Securities Information Processors ("SIPs").

Background

The SIPs currently each reside in a single location (Mahwah, New Jersey for SIAC, the processor for Tapes A and B securities; Carteret, New Jersey for Nasdaq UTP, the processor for Tape C securities). Each SIP receives quote and trade data from all NMS Plan Participants (i.e., exchanges and FINRA), consolidates this information into a national best bid and offer ("NBBO") and other calculated values (e.g. Limit Up/Limit Down ("LULD") price bands) for each security, then publishes the NBBO, the best quotes from each market, last-sale information, and regulatory messages to data recipients ("Recipients").

The SIP Operating Committee has improved system performance by directing the SIPs to develop lower latency and higher throughput processing (e.g., re-platforming with binary messaging, FPGA acceleration). The consolidation and re-broadcast process has become much faster over recent years, and is now typically accomplished in tens of microseconds instead of several milliseconds. As this processing time approaches zero, it becomes clear that the time required for trade and quote data to travel from Participant datacenter->SIP datacenter->Recipient datacenter, or "geographic latency", is a larger proportion of the total latency. Exhibit A, attached, illustrates this geographic latency.

Possible alternatives to the current model

There are a range of possible alternative data delivery models that can be considered to minimize time differences in when consolidated data and single market proprietary data is available in each major New Jersey-area datacenter. These alternatives include establishing instances of the existing single processors in more than one data center, replacing the single processor with many

competing consolidators, or eliminating the SIP altogether and shifting the burden of data consolidation to market participants (a la European equities markets).

In our view, both the competing consolidator and European models of data consolidation would introduce levels of complexity that would undermine the purpose of Regulation NMS to keep costs down for retail investors. Rather, we believe that a "Distributed SIP" implementation of the existing processors, as described below, would be the simplest of the alternative models and the most likely to continue to achieve the goals of Regulation NMS. Any alternative to the current model would, of course, need to be thoroughly evaluated by the Commission.

The "Distributed SIP" concept

In this Distributed SIP model: (1) the existing SIPs would establish instances of their systems in multiple data centers, preliminarily in each of Mahwah, Carteret, Secaucus, and Chicago; (2) Participants would publish their quote and trade data to each SIP instance; and (3) Recipients could consume consolidated data for Tapes A, B, and C securities from one or more of the SIP locations. As a result, market data would never have to travel more than one "hop" before it would be available to a Recipient in one of the four major datacenters (see examples below).

Each instance of the SIP would operate in parallel, consuming quote and trade data directly from the Participants, independently consolidating the data and republishing it to Recipients. Recipients would have the ability to connect to SIP data streams at one or more locations (commercial policy TBD). In this model, a recipient that wished to continue using the existing SIP sources (Tapes A and B via Mahwah, Tape C via Carteret) without making any changes would be free to do so, but could also choose to use one of the proposed new SIP sources (e.g., Secaucus or Carteret for Tape A).

Implementation considerations

Implementing a Distributed SIP system would reduce geographic latency, but would also introduce new complexities and require the Commission to address certain legal and policy issues. The following is a preliminary list of those new complexities and issues for consideration:

- Regulatory Treatment of Multiple NBBOs. Could a data recipient rely on the NBBO from any one of the four locations for Reg NMS regulatory compliance purposes? While each SIP instance would be expected to receive all updates from every exchange, those updates would arrive at each location at different points in time. As a result, the NBBO 'snapshot' for a specific security at any given point in time may be slightly different at the Mahwah SIP versus the Carteret SIP versus the Secaucus SIP.

  The Commission would need to provide regulatory assurance that market participants could continue to use one or more quote sources, including a blend of SIP or proprietary data feeds, for Rule 611 regulatory compliance. Market centers and broker-dealers would need to continue to have consistent procedures for arbitrating between data sources.

- Impact on Regulatory Messages. Regulatory messages currently calculated by the SIP, such as the LULD price bands, would also be impacted by the different message arrival sequence at each SIP location. To limit complexity, regulatory messages could be calculated by only one SIP instance (located in the same datacenter as a security's listing exchange) and then re-distributed to the other instances for re-publishing. This would reduce the risk of persistent inconsistencies in regulatory information from the disparate SIP instances.

- Disaster Recovery. Another implementation consideration is whether Chicago should continue to serve as a disaster recovery site, or if Recipients could subscribe to multiple New Jersey metro area instances as a back-up. If regulators want Chicago to remain a viable disaster recovery option to provide geographic diversity, Chicago instances of SIAC and Nasdaq UTP SIP could operate as a 4th Distributed SIP instance.

- Regulation NMS Rule 603. The Commission would need to consider changes to its existing rules, including Rule 603. For example, we believe it would be necessary for policymakers to evaluate whether a Distributed SIP model would be consistent with, or whether changes would be required to, the Rule 603(a)(1) requirement that an exclusive processor distribute information "on terms that are fair and reasonable" and the Rule 603(b) requirement that SIP Plans "provide for the dissemination of all consolidated information for an individual NMS stock through a single plan processor."[1] Similarly, the Commission would need to evaluate whether changes were required to Rule 603(c)(1) that "[n]o securities information processor. . . shall provide, in a context in which a trading or order-routing decision can be implemented, a display of any information with respect to quotations for or transactions in an NMS stock without also providing, in an equivalent manner, a consolidated display for such stock." Would four simultaneous transmissions be considered "consolidated" if they differ because of geographic latency?

Why the Distributed SIP concept?

There are several reasons the Distributed SIP concept may be more readily achieved than alternative re-design proposals.

1. Reg NMS Rule 603(b) requires that the SIP Plans "shall provide for the dissemination of all consolidated information for an individual NMS stock through a *single plan processor*." (emphasis added) Arguably because SIAC and Nasdaq UTP would remain the sole processors (albeit with multiple installations), this change could be implemented without revision to Reg NMS.

---

[1] *See* Regulation NMS Adopting Release, Securities Exchange Act Release No. 51808, 70 FR 37496, at 37569 (June 29, 2005) (File No. S7-10-04): "As a result, information users, particularly retail investors, will be able to obtain data from a single source that reflects the best quotations and most recent trade price for a security, no matter where such quotations and trade are displayed in the NMS."

2. Recipients would not be required to make any changes to their existing technology or procedures unless they choose to do so. Firms without sub-millisecond latency concerns could simply continue to consume Tapes A and B data from SIAC in Mahwah and Tape C data from Nasdaq UTP in Carteret, while those with geographic latency concerns could shift their consumption to the nearest possible instance of each SIP.
3. While the SIP outputs would differ between datacenters based on time and sequence of trade and quote messages, the systems' processing logic would be identical and provide consistent behavior.
4. Preserves the Operating Committee's clear responsibility for operating (through its processors) and administering a highly reliable Reg SCI system, versus the potentially ambiguous liabilities that may exist among many competing technology providers.

Recommendation

NYSE recommends the Commission undertake an analysis of the costs and benefits to the industry (institutional brokers, ATSs, retail investors, third-party data aggregators, institutional buy side, and exchanges) of de-centralizing the SIP architecture, either through multiple instances of the existing processors or an alternative method. The benefits would include some objective efficiencies, as well as the less tangible but equally important elimination of the perception that the SIP is structurally inferior to proprietary data products. The costs to the industry of any change to the SIP plant will be significant, and include technology changes, compliance changes, policy and procedure changes, regulatory surveillance changes, and other to be enumerated activities.

Respectfully submitted,

Michael Blaugrund

Exhibit A: Examples of Geographic Latency

Example 1: Tape B quote update on BZX, with Recipient in Secaucus

Cboe BZX operates in the Secaucus, New Jersey datacenter. SIP Recipients located in Secaucus currently receive a Tape B quote update on Cboe BZX only after it has travelled to Mahwah, been processed by SIAC, and returned to Secaucus. This takes ~480 microseconds, ~400 of which are due to travel time.

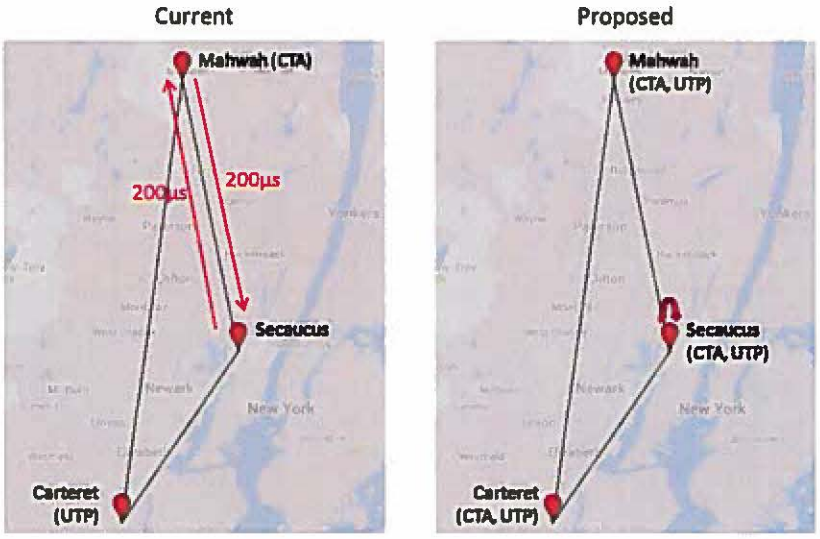

If instead a Distributed SIAC instance were located in Secaucus, the market data update would be subject only to the 80 microsecond processing time--an 83% speed improvement.

Example 2: Tape C quote update on Arca, with Recipient in Mahwah

Because the Mahwah, New Jersey datacenter does not host a SIP for Tape C securities, a recipient of SIP data in Mahwah would receive a Tape C quote or trade on Arca only after it has travelled to Carteret, been processed by the Nasdaq UTP SIP, and returned to Mahwah. This takes ~720 microseconds, ~700 of which are due to travel time.

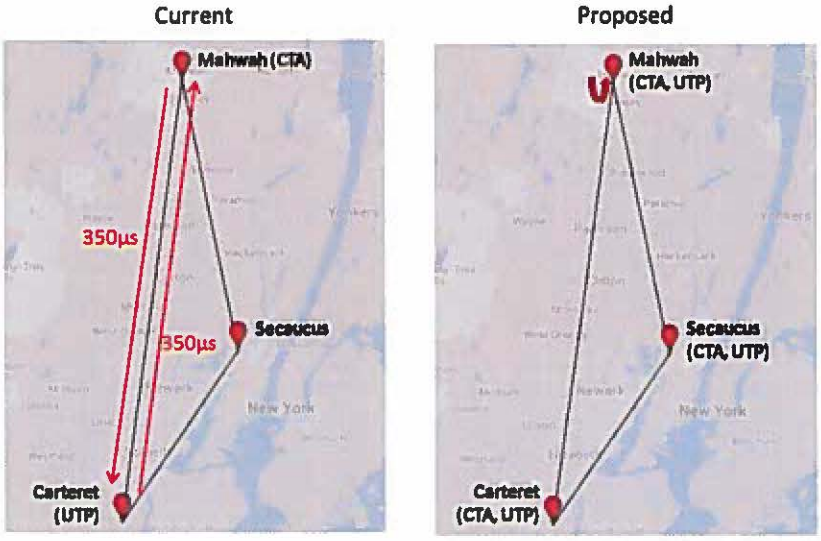

If instead a UTP SIP instance were located in Mahwah, the market data update would be subject only to the 20 microsecond processing time--a 90% speed improvement.

Example 3: Tape C quote update on Arca, with Recipient in Secaucus

Consider the same quote update as above, but for a recipient in Secaucus, New Jersey. Currently, the Recipient would receive the update after it has travelled to Carteret, New Jersey, been processed by the Nasdaq SIP, and then traveled to Secaucus. This takes ~520 microseconds, ~500 of which are due to travel time.

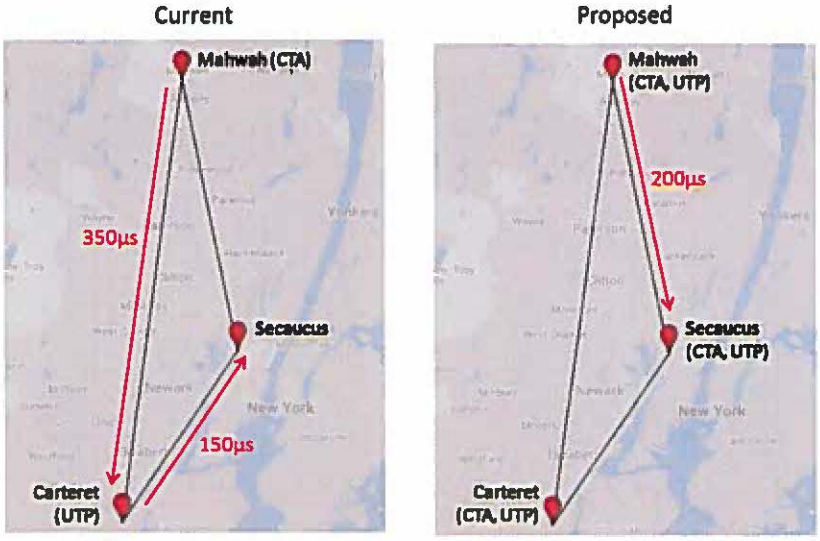

If instead a Nasdaq SIP instance were located in Secaucus, the market data update would be sent directly from Mahwah to Secaucus (200 microseconds) and then be subject to UTP SIP processing (20 microseconds) --a 58% speed improvement.